

ESTIMATION AND PREDICTION IN SPARSE AND UNBALANCED TABLES

QINGYUAN ZHAO, TREVOR HASTIE AND DARYL PREGIBON

University of Pennsylvania, Stanford University and Google Inc.

ABSTRACT. We consider the problem where we have a multi-way table of means, indexed by several factors, where each factor can have a large number of levels. The entry in each cell is the mean of some response, averaged over the observations falling into that cell. Some cells may be very sparsely populated, and in extreme cases, not populated at all. We might still like to estimate an expected response in such cells. We propose here a novel hierarchical ANOVA (HANOVA) representation for such data. Sparse cells will lean more on the lower-order interaction model for the data. These in turn could have components that are poorly represented in the data, in which case they rely on yet lower-order models. Our approach leads to a simple hierarchical algorithm, requiring repeated calculations of sub-table means of modified counts. The algorithm has shown superiority over the unshrunk methods in both simulations and real data sets.

1. INTRODUCTION

Prediction with factorial features has been studied for a long time in statistics. Recently there are many arising applications of this kind, but with much larger data size than before. For instance, consider data on restaurant ratings, as provided by Zagat, Yelp or other such services. In addition, the restaurants can usually be classified according to various factors, such as zip-code or geographical region (at least 10000 levels), type of cuisine (Italian, French, etc, potentially dozens of levels), and price category (often 1-5 stars). One important task is to estimate the average rating for a particular kind of restaurants defined by these factors, which can be used to answer questions like “What is the most popular cuisine in San Francisco?” or “Do expensive restaurants usually get better ratings?”.

Such data sets and related questions have been studied thoroughly by statisticians from the first formal proposal of analysis of variance (ANOVA) in Fisher (1918). In most of the analysis of variance literature, the focus is how to estimate the relative importance of different variance components and how to test for their statistical significance. The problem studied in this paper differs from classical ANOVA in the following aspects:

- The focus is estimating and predicting cell means, whereas the main objective of ANOVA is to find important factors or interactions and test for their significance.
- The observed data table is sparse and highly unbalanced. In many applications, we may only have a small proportion of the cells observed with quite different weights. For example, one zip code may have no Ethiopian restaurant, only one unpopular Japanese restaurant with a few dozen ratings, but many American restaurants with thousands of ratings in total.
- The size of the data is usually very large and we also have a large number of main effects and interactions to estimate, so it is very important to find an computationally efficient algorithm while maintaining good statistical properties.

E-mail address: qyzhao@wharton.upenn.edu.

Key words and phrases. factorial data, hierarchical modeling, ANOVA.

Qingyuan Zhao was a Ph.D. candidate at Stanford University when this work completed. Trevor Hastie was partially supported by grant DMS-1407548 from the National Science Foundation, grant 5R01 EB 001988-21 from the National Institutes of Health, and a Google research grant.

In the rest of the paper, we first illustrate our method with the restaurant ratings example in Section 2, then describe the statistical problem we are trying to solve and some previous literature in Section 3. In Section 4.1 we prove some theoretical properties of our algorithm when the observed data is balanced. This gives us an empirical choice of penalty parameter, as described in Section 4.2. The rest of the paper is devoted to algorithm implementation (Section 5) and simulations and some real data results (Section 6).

2. ESTIMATION IN THREE-WAY TABLES

Our proposed method is general can be applied to any number of factors, but this would require a somewhat technical representation. In this section we will restrict ourselves to three factors (F_1 , F_2 and F_3) to demonstrate the key elements in the algorithm.

Suppose these factors have I , J , and K levels respectively. The restaurant ratings are a collection of n_{ijk} measurements $y_{ijk\ell}$, $\ell = 1, \dots, n_{ijk}$ at each cell (i, j, k) , and we are interested in predicting the cell mean $\mu_{i,j,k} = E(Y|F_1 = i, F_2 = j, F_3 = k)$. We summarize our data by the cell means $\bar{y}_{ijk} = y_{ijk\cdot}/n_{ijk}$ and a weight n_{ijk} . For simplicity, we drop the bar, and refer to our aggregate data as y_{ijk} .

For some cells n_{ijk} could be very small, even zero, but we would still like to have a reasonable estimate. In this case we would like to shrink our estimate towards a more stable number that borrows strength from similar cells. This suggests a Bayesian mixed-effects framework Diggle et al. (1994), where we assume the μ_{ijk} are random, say Gaussian, with distribution $\mu_{ijk} \sim N(\gamma_{ijk}, \sigma_\mu^2)$. For the moment assume γ_{ijk} is known. If we assume the original measurements $y_{ijk\ell}|\mu_{ijk} \sim N(\mu_{ijk}, \sigma^2)$ and are all independent, then the negative log-posterior likelihood for μ_{ijk} given data is proportional to

$$(2.1) \quad L(\mu) = \sum_{i,j,k} n_{ijk} (y_{ijk} - \mu_{ijk})^2 + \lambda \sum_{i,j,k} (\mu_{ijk} - \gamma_{ijk})^2,$$

where $\lambda = \sigma^2/\sigma_\mu^2$. The posterior mode is simple to characterize:

$$(2.2) \quad E(\mu_{ijk}|\mathbf{y}) = \frac{n_{ijk}y_{ijk}}{n_{ijk} + \lambda} + \frac{\lambda\gamma_{ijk}}{n_{ijk} + \lambda},$$

a simple weighted average of the observed mean y_{ijk} and γ_{ijk} , with more emphasis on the former when n_{ijk} is large. Now γ_{ijk} is not known, and so we can represent it by a parametric model $\gamma_{ijk}(\theta)$. Empirical Bayes amounts to estimating θ by maximizing the marginal likelihood (integrating out μ), and then using $\gamma_{ijk}(\hat{\theta})$ as the target of shrinkage in (2.2).

Barry (1990), in a balanced two-way layout, proposed a simple main-effects parametrization $\gamma_{ij} = \alpha^0 + \alpha_i^1 + \alpha_j^2$, with $\sum_i \alpha_i^1 = 0$, and $\sum_j \alpha_j^2 = 0$. One can show in this case that the empirical Bayes procedure amounts to maximizing the following likelihood

$$(2.3) \quad L(\mu) = \sum_{i,j} (y_{ij} - \mu_{ij})^2 + \lambda \sum_{i,j} (\mu_{ij} - \mu_{i\cdot} - \mu_{\cdot j} + \mu_{\cdot\cdot})^2,$$

where $\mu_{i\cdot} = \sum_{j=1}^J \mu_{ij}/I$ etc. Now one can show two things:

- We can write (2.3) in vector notation as

$$(2.4) \quad L(\mu) = (\mathbf{y} - \mu)^T(\mathbf{y} - \mu) + \lambda \mu^T(\mathbf{I} - \mathbf{P}_A)\mu,$$

where \mathbf{P}_A is the main-effects ANOVA projection operator. This has solution

$$(2.5) \quad \hat{\mu} = (\mathbf{I} + \lambda(\mathbf{I} - \mathbf{P}_A))^{-1}\mathbf{y}.$$

- This solution can be shown to have the much simpler form

$$(2.6) \quad \hat{\mu} = \frac{\mathbf{y}}{1 + \lambda} + \frac{\lambda \tilde{\mathbf{y}}}{1 + \lambda},$$

where $\tilde{\mathbf{y}} = \mathbf{P}_A \mathbf{y}$, which is the main-effects ANOVA fit. This is of the same form as (2.2), with $n_{ij} = 1$ and $\gamma_{ij} = \tilde{y}_{ij}$

This nice simplification disappears if we include weights in (2.3) or (2.4), although we do get a closed form expression for the solution along the lines of (2.6) (but involving weights). What we like about (2.6), apart from its simplicity, is that it is easy to compute. The ANOVA fit \tilde{y}_{ij} requires simple marginal means of y_{ij} along the two factors, and the overall mean. There are several problems though. This simplicity, both in representation and especially in computation, goes away when each observation has weights. In addition, it may be that the main-effects ANOVA model is not well estimated for some values of i or j , because of sparsity in these margins.

We propose a method that has the simplicity of (2.2) for a weighted model, that has a multi-level hierarchical structure, and is easy to compute.

Hierarchical penalized ANOVA model for a three-way table

0. Fit $\mu_{ijk}^{(0)} = \bar{y} \dots$, the overall weighted mean of all the y s.
1. Fit the additive model $\mu_{ijk}^{(1)} = \alpha_i^1 + \alpha_j^2 + \alpha_k^3$ by solving the following weighted penalized least squares (WPLS) problem:

$$\min_{\alpha} \sum_{i,j,k} n_{ijk} (y_{ijk} - (\alpha_i^1 + \alpha_j^2 + \alpha_k^3))^2 + \lambda_1 \cdot \sum_{n_{ijk} > 0} (\alpha_i^1 + \alpha_j^2 + \alpha_k^3 - \mu_{ijk}^{(0)})^2.$$

This can be solved using a simple backfitting algorithm Hastie and Tibshirani (1990), where each step is a version of (2.2) that only requires table summing. Here we shrink to the overall mean, so if a particular one-way marginal count is low, the shrinkage will still kick in.

2. Fit the second-order interaction model $\mu_{ijk}^{(2)} = \psi_{ij}^{12} + \psi_{jk}^{23} + \psi_{ik}^{13}$ by solving the WPLS problem:

$$\min_{\psi} \sum_{i,j,k} n_{ijk} (y_{ijk} - (\psi_{ij}^{12} + \psi_{jk}^{23} + \psi_{ik}^{13}))^2 + \lambda_2 \cdot \sum_{n_{ijk} > 0} (\psi_{ij}^{12} + \psi_{jk}^{23} + \psi_{ik}^{13} - \mu_{ijk}^{(1)})^2.$$

This is the case we were concerned about in Section 1, when some two-way tables have sparse counts. In this case, that term is shrunk more towards its parent.

3. Fit the full third-order model $\mu_{ijk}^{(3)}$ by solving the WPLS problem

$$\min_{\mu} \sum_{i,j,k} n_{ijk} (y_{ijk} - \mu_{ijk})^2 + \lambda_3 \sum_{n_{ijk} > 0} (\mu_{ijk} - \mu_{ijk}^{(2)})^2.$$

This is where we started, i.e. model (2.2), with $\gamma_{ijk} = \tilde{y}_{ijk}^{(2)}$. The term we shrink to has been regularized to accommodate lower-level sparsity.

We have omitted the details of backfitting in steps 1 and 2. Consider using coordinate-descent to solve (2), where for example we hold α_j^2 and α_k^3 fixed at their current estimates and solve for α_i^1 . By collecting terms, this is again of the form (2.2), and so the update involves simple table averaging of a modified response and shrinkage target.

There are many interesting details that haven't been discussed yet. For example, we have the same λ at each level; these could be different. We can take the empirical Bayes analogy further and use variance components to suggest values for λ , or at least *relative values*. We will discuss this in Section 4. Another option is to terminate the hierarchy early. In the restaurant ratings example, three way table is manageable even with many levels for each factor. However, when the number of factors

grows, the number of cells grows at an exponential rate. In this case we might truncate the hierarchy early, e.g. at step 2, and use $\tilde{y}_{ijk}^{(2)} = \tilde{\psi}_{ij}^{12} + \tilde{\psi}_{jk}^{23} + \tilde{\psi}_{ik}^{13}$ as the estimate. Note in this case we would only need to store and compute at the level of two-way tables.

3. EXTENSION TO GENERAL TABLES

Let's extend our algorithm to a general table with $m \geq 1$ factors. All the cell means are represented by a vector \mathbf{y} . Each cell mean is indexed by its factor levels $I = (i_1, i_2, \dots, i_m)$. For example, if $m = 3$, then the cell that is in the 2nd, 4th and 1st level of the three factors corresponds to $y_I = y_{241}$. Every cell is also associated with a weight, denoted by n_I .

Let Ω be the set of observed indices. The general statistical question we are trying to address is: suppose the ratings follow an underlying statistical model $\mathbf{y} = \boldsymbol{\mu} + \boldsymbol{\epsilon}$ and $\mu_I = f(I; \boldsymbol{\beta})$ for some deterministic or random function f , if we have only observed \mathbf{y}_Ω , how do we estimate the underlying cell means $\boldsymbol{\mu}$?

3.1. General Algorithm. In this paper, our algorithm assumes a hierarchy of models f_0, f_1, \dots, f_m increasing in complexity. Each time the algorithm uses fits from the previous model f_k as the prior for a new and more complex model f_{k+1} . This is a very general hierarchical model and we may choose different sequence of model functions in different applications.

In our problem, the model function f_k , $0 \leq k \leq m$ is a k -th order interaction model, i.e.

$$(3.1) \quad \mu_I^{(k)} = f_k(I; \boldsymbol{\beta}^{(k)}) = \sum_{J \subset [m], |J|=k} \beta_{I_J}^J$$

Here J is a subset of $[m] = \{1, 2, \dots, m\}$, indicating the factors we are considering; $\beta_{I_J}^J$ is the "mean" effect of cells that are in levels I_J of the factors specified in J , and $\boldsymbol{\beta}^{(k)} = \{\beta^J : J \subset [m], |J| = k\}$ is the model parameter for f_k .

Our algorithm is summarized in Algorithm 1. Notice that this idea can be generalized to solve any hierarchical model, not necessarily the full additive model in (3.1).

Algorithm 1 Generic algorithm to solve (3.1)

for $k = 1, \dots, m + 1$ **do**

 Estimate $\boldsymbol{\beta}^{(k)}$ by $\hat{\boldsymbol{\beta}}^{(k)}$, the Bayes estimate given \mathbf{y} (data) and $\boldsymbol{\beta}^{(k-1)} = \hat{\boldsymbol{\beta}}^{(k-1)}$ (prior).

end for

3.2. Previous Methods. Many statisticians studied this kind of problem in the 1980's using Bayes models. One typical model (DuMouchel and Harris, 1983) is

$$(3.2) \quad \begin{aligned} y_{ij} &= \theta_{ij} + \epsilon_{ij}, \\ \theta_{ij} &= \mu + \alpha_i + \beta_j + \delta_{ij}, \end{aligned}$$

and they also want to extrapolate for some missing cells in the table. The empirical Bayes method is to put a normal prior on α and β

$$(3.3) \quad \alpha, \beta \sim N(0, V)$$

and estimate the hyperprior variance parameter V using the data. A more general Bayesian ANOVA model and many more examples are considered in the tutorial paper by Casella (1992). See also Searle et al. (1992, ch. 9) and Searle (2006).

More recently, Beran (2005) considered ANOVA from a pure shrinkage and optimization view. The author considers the whole class of estimators that solve penalized least squares, which may correspond to some Bayes modeling or not. Then among all the possible shrinkage estimators, the paper tries to find the one that minimize estimated risk. This is a very sound approach from a theoretical point of view, but the algorithm proposed in the paper scales poorly and couldn't be used in real data sets.

In Gelman (2005), the author uses a hierarchical Bayes model to analyze Analysis of Variance in an universal way. The statistical model is

$$(3.4) \quad y_i = \sum_{m=0}^M \beta_{j_i}^{(m)},$$

where m stands for a batch of regression coefficients (or in terms of ANOVA, variance component). Then he assumes normal priors on these coefficients

$$(3.5) \quad \beta_j^{(m)} \sim N(0, \sigma_m^2), \quad \forall j = 1, \dots, J_m, \quad m = 1, \dots, M$$

and puts hierarchical prior on σ_m^2 . This bears some similarity to our statistical model described in 3.1. Gelman's paper focuses mostly on the relationship between this hierarchical Bayes model with classical ANOVA and how to use Gibbs sampler to sample from the posterior distribution of σ_m^2 .

Following this idea, a more recent paper by Volfovsky and Hoff (2014) considers a three-way factorial model

$$(3.6) \quad y_{ijkl} = \mu + a_i + b_j + c_k + (ab)_{ij} + (ac)_{ik} + (bc)_{jk} + (abc)_{ijk} + \epsilon_{ijkl}$$

where ϵ_{ijkl} are i.i.d. $N(0, \sigma^2)$. However, the effects here are possibly correlated, which makes the problem different from the usual settings. The authors put a normal prior with zero mean and general covariance matrix on the effects and treat the covariance matrices as parameters with invert-Wishart priors. By plugging in empirical Bayes estimate for the hyperpriors, they can run a Gibbs sampler to find the posterior.

However, despite the flexibility of hierarchical Bayes model, it usually needs a Markov Chain Monte Carlo (MCMC) algorithm to sample from posterior distribution. This is unacceptable for more than a few hundred cells. Also, as noticed in Section 1, our focus is on estimating and predicting cell means instead of inference about variance components.

From the perspective of high dimensional estimation, shrinkage is known to be very effective, for example the James-Stein estimator dominates maximum likelihood estimator (Efron, 2010). Recently, in a multi-task averaging problem, Feldman et al. (2012) uses James-Stein estimator with empirically estimated covariance matrix and claims to outperform James-Stein in some occasions. However, the multi-task averaging problem has no associated covariates, and the structural information of these factors are crucial in our algorithm.

In light of the hierarchical Bayes model considered so far, our algorithm can be viewed as an empirical solution to a hierarchical Bayes model.

4. THEORY

In this section we develop some theoretical results for HANOVA. These also provide a method to empirically select the regularization parameter λ in the algorithm.

4.1. Balanced Table. Let's first assume all the observed cells have the same weight (i.e. same number of observations).¹ Assume we observe n cells and the responses are centered. At each stage, we are maximizing the log "likelihood" function

$$(4.1) \quad -\frac{1}{2} [\|\mathbf{y} - \boldsymbol{\mu}^{(k)}\|^2 + \lambda_k \|\boldsymbol{\mu}^{(k)} - \boldsymbol{\mu}^{(k-1)}\|^2]$$

subject to $\boldsymbol{\mu}^{(k)} \in \mathcal{S}_k$. Here \mathcal{S}_k is the linear subspace of \mathbb{R}^n generated by all possible k -th order effects. This implies the subspaces are nested $\mathcal{S}_1 \subset \mathcal{S}_2 \subset \dots \subset \mathcal{S}_k$. Say the dimension of \mathcal{S}_k is d_k and the projection matrix onto \mathcal{S}_k is \mathbf{P}_k . This means $\mathbf{P}_k = \mathbf{U}_k \mathbf{U}_k^T$ where the columns of $\mathbf{U}_k \in \mathbb{R}^{n \times d_k}$ are orthonormal basis for \mathcal{S}_k , i.e. $\mathbf{U}_k^T \mathbf{U}_k = \mathbf{I}_{d_k}$. Because the subspaces are nested, we

¹This is different from what a "balanced table" is meant in most of the statistics literature. In our paper, we say a table is balanced if all the observed cells have the same number of observations, but the table could contain many empty cells.

can assume the first d_{k-1} columns of \mathbf{U}_k are \mathbf{U}_{k-1} , i.e. $\mathbf{U}_k = \begin{pmatrix} \mathbf{U}_{k-1} & \mathbf{V}_k \end{pmatrix}$. Moreover suppose $\mathbf{U} = \begin{pmatrix} \mathbf{U}_K & \mathbf{V} \end{pmatrix} \in \mathbb{R}^{n \times n}$ is an orthogonal matrix.

Consider a hierarchical Bayes model that is a special case of (3.1).

$$\begin{aligned}
 \mathbf{y}|\boldsymbol{\mu} &\sim N(\boldsymbol{\mu}, \sigma^2 \mathbf{I}_n), \sigma^2 \text{ is known} \\
 \boldsymbol{\mu}|\boldsymbol{\beta}^{(m)} &\sim N(\mathbf{U}_m \boldsymbol{\beta}^{(m)}, \sigma_m^2 \mathbf{I}_n), \\
 \boldsymbol{\beta}^{(m)}|\boldsymbol{\beta}^{(m-1)} &\sim N(\mathbf{U}_m^T \boldsymbol{\mu}^{(m-1)}, \sigma_{m-1}^2 \mathbf{I}_{d_m}), \boldsymbol{\mu}^{(m)} = \mathbf{U}_m \boldsymbol{\beta}^{(m)} \\
 &\vdots \\
 \boldsymbol{\beta}^{(k)}|\boldsymbol{\beta}^{(k-1)} &\sim N(\mathbf{U}_k^T \boldsymbol{\mu}^{(k-1)}, \sigma_{k-1}^2 \mathbf{I}_{d_k}), \boldsymbol{\mu}^{(k)} = \mathbf{U}_k \boldsymbol{\beta}^{(k)} \\
 &\vdots \\
 \boldsymbol{\beta}^{(1)}|\boldsymbol{\beta}^{(0)} &\sim N(\mathbf{U}_1^T \boldsymbol{\mu}^{(0)}, \sigma_0^2 \mathbf{I}_{d_1}), \boldsymbol{\mu}^{(1)} = \mathbf{U}_1 \boldsymbol{\beta}^{(1)}.
 \end{aligned} \tag{4.2}$$

As the next theorem indicates, this hierarchical model is closely related to maximizing the log-likelihood (4.1). In the usual ANOVA model, each factor may have totally different main effects (or interactions) and we can test for the significance of them. However, as pointed out in Section 1, we are no longer interested in any individual factor. By using orthogonal matrices in the hierarchical model (4.2), we implicitly treat the effects and interactions of the observable factors as "randomly chosen" directions in the column space of \mathbf{U}_m .

Theorem 1. *For balanced table, each step of the hierarchical procedure (Algorithm 1) for (4.2) is equivalent to maximizing (4.1) subject to $\boldsymbol{\mu}^{(k)} \in \mathcal{S}_k$, with*

$$\lambda_k = \frac{\sigma^2 + \sigma_m^2 + \dots + \sigma_k^2}{\sigma_{k-1}^2} \tag{4.3}$$

and $\boldsymbol{\beta}^{(k-1)}$ being the fit from last step.

Proof. See Appendix A.1. □

Theorem 1 also indicates we can estimate all the variance parameters σ_k^2 from the data and then compute a vector of $\boldsymbol{\lambda}$. Let's assume the ratings are centered (by the global mean) so we put $\boldsymbol{\beta}^{(0)} = \mathbf{0}$. Then the multi-level hierarchical model (4.2) is actually a linear random effects model

$$\begin{aligned}
 \mathbf{y}|\boldsymbol{\mu} &\sim N(\boldsymbol{\mu}, \sigma^2 \mathbf{I}_n), \sigma^2 \text{ is known} \\
 \boldsymbol{\mu}|\boldsymbol{\beta}^{(m)} &\sim N(\mathbf{U}_m \boldsymbol{\beta}^{(m)}, \sigma_m^2 \mathbf{I}_n), \\
 \boldsymbol{\beta}^{(m)} &\sim N(\mathbf{0}, \boldsymbol{\Sigma}), \\
 \boldsymbol{\Sigma} &= \begin{pmatrix} \tau_0^2 \mathbf{I}_{d_1} & & & \\ & \tau_1^2 \mathbf{I}_{d_2-d_1} & & \\ & & \ddots & \\ & & & \tau_{m-1}^2 \mathbf{I}_{d_m-d_{m-1}} \end{pmatrix}, \\
 \tau_i^2 &= (\sigma_i^2 + \dots + \sigma_{m-1}^2), \quad i = 0, \dots, m.
 \end{aligned} \tag{4.4}$$

Now we can state our main theorem that guarantees the effectiveness of our algorithm for balanced table.

Theorem 2. *For balanced table, if we use (4.3) to compute $\boldsymbol{\lambda}$, the solution to our algorithm is the posterior mean of $\boldsymbol{\beta}$ in the linear random effects model (4.4).*

Proof. See Appendix A.2. □

In reality we don't know what $\sigma_0^2, \sigma_1^2, \dots, \sigma_m^2$ are, but we can use various methods developed in ANOVA to estimate them. By doing this, our algorithm is equivalent to an empirical Bayes solution of the multi-level hierarchical model.

Following the general principle of variance component analysis, we can compute the expectation of certain quadratic forms of \mathbf{y} , then use these quadratic forms to find an unbiased estimator of σ_k^2 . For example,

$$\begin{aligned}
 \mathbb{E}[\mathbf{y}^T \mathbf{U}_k \mathbf{U}_k^T \mathbf{y}] &= \text{tr}[\mathbf{U}_k \mathbf{U}_k^T \text{Var}(\mathbf{y})] \\
 &= \text{tr}[\mathbf{U}_k \mathbf{U}_k^T (\sigma^2 \mathbf{I}_n + \sigma_m^2 \mathbf{I}_n + \mathbf{U}_m \Sigma \mathbf{U}_m^T)] \\
 &= d_k(\sigma^2 + \sigma_m^2) + \sum_{j=0}^{k-1} (d_{j+1} - d_j) \tau_j^2, \\
 &\quad \forall k = 1, \dots, m
 \end{aligned}
 \tag{4.5}$$

Also

$$\begin{aligned}
 \mathbb{E}[\mathbf{y}^T \mathbf{U} \mathbf{U}^T \mathbf{y}] &= \mathbb{E}[\|\mathbf{y}\|^2] \\
 &= \text{tr}[(\sigma^2 + \sigma_m^2) \mathbf{I}_n + \mathbf{U}_m \Sigma \mathbf{U}_m^T] \\
 &= n(\sigma^2 + \sigma_m^2) + \sum_{j=0}^{m-1} d_{j+1} \sigma_j^2
 \end{aligned}
 \tag{4.6}$$

Thus we have $m + 1$ equations and $m + 1$ parameters (recall we assume σ^2 is known), the unbiased estimator of the prior variances can be obtained by solving this linear system.

In fact, if we call $\mathbf{U}_{m+1} = \mathbf{U}$ and $d_{m+1} = n$, we have

$$\begin{aligned}
 \mathbb{E}[\mathbf{y}^T \mathbf{U}_{k+1} \mathbf{U}_{k+1}^T \mathbf{y} - \mathbf{y}^T \mathbf{U}_k \mathbf{U}_k^T \mathbf{y}] &= (d_{k+1} - d_k) \tau_k^2, \\
 &\quad \forall k = 0, \dots, m
 \end{aligned}
 \tag{4.7}$$

So

$$\begin{aligned}
 \hat{\sigma}_k^2 &= \frac{\mathbf{y}^T \mathbf{U}_{k+1} \mathbf{U}_{k+1}^T \mathbf{y} - \mathbf{y}^T \mathbf{U}_k \mathbf{U}_k^T \mathbf{y}}{d_{k+1} - d_k} \\
 &\quad - \frac{\mathbf{y}^T \mathbf{U}_{k+2} \mathbf{U}_{k+2}^T \mathbf{y} - \mathbf{y}^T \mathbf{U}_{k+1} \mathbf{U}_{k+1}^T \mathbf{y}}{d_{k+2} - d_{k+1}}, \\
 &\quad k = 0, \dots, m-1 \\
 \hat{\sigma}_m^2 &= \frac{\mathbf{y}^T \mathbf{y} - \mathbf{y}^T \mathbf{U}_m \mathbf{U}_m^T \mathbf{y}}{n - d_m} - \sigma^2
 \end{aligned}
 \tag{4.8}$$

are unbiased for estimating σ_k^2 . Note that it is also possible to treat σ^2 as a tuning parameter, but usually σ^2 can be estimated quite accurately from the data.

The above formulae is only one specific (perhaps the easiest) choice of estimating equations. More general methods in analysis of variance for unbalanced tables, such as Restricted Maximum Likelihood (REML) can also be used (see, for example, Searle et al. (1992)).

4.2. Unbalanced Table. If the observed table is unbalanced, then we cannot prove any exact conclusions like 1 or 2. In this case, the log likelihood function we maximize is

$$-\frac{1}{2}[(\mathbf{y} - \boldsymbol{\mu}^{(k)})^T \mathbf{N}(\mathbf{y} - \boldsymbol{\mu}^{(k)}) + \lambda_k \|\boldsymbol{\mu}^{(k)} - \boldsymbol{\mu}^{(k-1)}\|^2]
 \tag{4.9}$$

where \mathbf{N} is a diagonal matrix with weights.

Similar to the previous section, we may sequentially maximize (4.9) for $k = 1, 2, \dots, m$ to estimate $\boldsymbol{\mu}$ and $\boldsymbol{\beta}$ in the following linear random effects model

$$(4.10) \quad \begin{aligned} \mathbf{y}|\boldsymbol{\mu} &\sim \mathcal{N}(\boldsymbol{\mu}, \sigma^2 \mathbf{N}^{-1}), \quad \sigma^2 \text{ is known} \\ \boldsymbol{\mu}|\boldsymbol{\beta}^{(m)} &\sim \mathcal{N}(\mathbf{U}_m \boldsymbol{\beta}^{(m)}, \sigma_m^2 \mathbf{I}_n), \\ \boldsymbol{\beta}^{(m)} &\sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}), \end{aligned}$$

Here $\boldsymbol{\Sigma}$ is the same as the one (4.4). This model is simply the unbalanced version of (4.4).

The theory we derived in Section 4.1 can be used to give us a vector of reasonable penalty parameters $\boldsymbol{\lambda}$. One way to do this is to pretend all the observed cells have the same weight and use equations in (4.8) to estimate all the σ_k^2 . Then one can compute ratios between the variance components (i.e. (4.3)) to get $\boldsymbol{\lambda}$. The estimate $\hat{\sigma}_k^2$ is still unbiased, because the estimation equations (4.5) and (4.6) still hold true after replacing σ^2 with $\sigma^2 \text{tr}(\mathbf{N}^{-1}/n)$.

5. IMPLEMENTATION

In this section we discuss some implementation details of HANOVA.

5.1. Preprocessing. In real applications, many data sets are not summarized in the format of cell means and weights. Usually we may have multiple units in a cell and every unit receive multiple reviews. A mixed-effect model describing this is

$$(5.1) \quad \begin{aligned} y_{ci} &\sim \mathcal{N}(\mu_c + \alpha_i, \sigma_r^2/n_i), \\ \alpha_i &\sim \mathcal{N}(0, \sigma_u^2). \end{aligned}$$

Here μ_c is the cell mean, α_i is the unit effect, σ_r^2 is the variance of user's rating. σ_r may actually depend on the unit, and the procedure below can be slightly modified to this heterogeneous case. n_i is the number of reviewers of that restaurant, σ_u^2 is the variance of the unit effect.

The preprocessing procedure for this model is described in Algorithm 2. Note that since we use the estimated variance directly as weight in step 3, the σ^2 defined in (4.2) is 1.

Algorithm 2 Data Preprocessing for HANOVA

- 1: Estimate σ_u^2 and σ_r^2 based on all the observed units.
- 2: Plug $\hat{\sigma}_u^2$ and $\hat{\sigma}_r^2$ in (5.1) and estimate the μ_c by

$$\hat{\mu}_c = \frac{\sum_i \frac{y_{ci}}{\hat{\sigma}_u^2 + \hat{\sigma}_r^2/n_i}}{\sum_i \frac{1}{\hat{\sigma}_u^2 + \hat{\sigma}_r^2/n_i}} \sim \mathcal{N}(\mu_c, \frac{1}{\sum_i \frac{1}{\hat{\sigma}_u^2 + \hat{\sigma}_r^2/n_i}})$$

- 3: Use $\hat{\mu}_c$ and $\sum_i \frac{1}{\hat{\sigma}_u^2 + \hat{\sigma}_r^2/n_i}$ as value and weight in HANOVA to obtain regularized estimate of μ_c , denote by $\hat{\hat{\mu}}_c$.
- 4: Estimate individual restaurant effect α_i by shrinking y_{ci} towards $\hat{\hat{\mu}}_c$, using the following formula

$$\hat{\hat{\mu}} + \hat{\alpha}_i = \frac{n_i y_{ci} / \hat{\sigma}_r^2 + \hat{\hat{\mu}}_c / \hat{\sigma}_u^2}{n_i / \hat{\sigma}_r^2 + 1 / \hat{\sigma}_u^2}$$

5.2. Normal Equations. In our algorithm, the k -th order model is

$$\min_{\boldsymbol{\beta}^{(k)}} \sum_{I \in \Omega} n_I (y_I - \mu_I^{(k)})^2 + \lambda_k \sum_{I \in \Omega} (\mu_I^{(k)} - \mu_I^{(k-1)})^2$$

Here $\mu_I^{(k)} = \sum_{J \subset [m], |J|=k} \beta_{I,J}^J$. Differentiate the above loss with respect to each μ_L^J and equate to zero

$$\sum_{I_J=L} (n_I + \lambda_k) \sum_{K \subset [m], |K|=k} \beta_{I,K}^K = \sum_{I_J=L} (n_I y_I + \lambda_k \mu_I^{(k-1)}),$$

$$\forall J \subset [m], |J| = k, L.$$

So the coefficient of β_M^K in the (J, L) -th equation is

$$z_{K,M}^{J,L} = \begin{cases} 0 & , K = J, M \neq L \\ \sum_{I_J=L} (n_I + \lambda_k) & , K = J, M = L \\ \sum_{I_J=L, I_K=M} (n_I + \lambda_k) & , K \neq J \end{cases}$$

Now the problem reduces to solve a large system of linear equations

$$\sum_{K,M} z_{K,M}^{J,L} \beta_M^K = \sum_{I_J=L} (n_I y_I + \lambda_k \mu_I^{(k-1)}), \quad \forall J \subset [m], |J| = k, L.$$

5.3. Backfitting. We implemented the block coordinate descent algorithm (or backfitting algorithm) to solve the above equations. Each block contains β^J , i.e. all the effects for certain margins. For example, when J contains only one factor, this block contains all the main effects associated with that factor. The psuedocode of our algorithm is in Algorithm 3.

Note that the basic building block of every iteration is computing $\sum_{I_J=L} (n_I + \lambda_k) \mu_I^{(old)}$. For a fixed J , this amounts to a weighted table sum over J^C . This operation can be easily parallelized.

Algorithm 3 psuedo-code of HANOVA

```

1: Choose maxk from  $1, \dots, m$ .
2: for  $k = 1 \rightarrow \text{maxk}$  do
3:   for all  $|J| = k, L \in \mathcal{L}(J)$  do
4:     initialize  $\beta_L^J \leftarrow \frac{\bar{y}}{\binom{m}{k}}$ 
5:      $w_L^J \leftarrow \sum_{I_J=L} (n_I y_I + \lambda_k \mu_I^{(k-1)})$ 
6:      $z_{J,L}^{J,L} \leftarrow \sum_{I_J=L} (n_I + \lambda_k)$ 
7:   end for
8:    $\mu^{(k)} = \mu^{(k-1)}$ 
9:   repeat
10:    for all  $|J| = k$  do
11:       $\mathbf{s}^J \leftarrow \text{apply}(\mu^{(k)}(\mathbf{n} + \lambda_k), J^C, \text{sum})$ 
12:       $\beta^J \leftarrow \beta^J + (\mathbf{u}^J - \mathbf{s}^J) / \mathbf{z}_{J^c}^J$ 
13:      Update  $\mu^{(k)}$ 
14:    end for
15:  until  $\mu$  converges.
16: end for

```

6. RESULTS

6.1. Simulations. We first use simulations to verify the optimality of HANOVA. We simulate 200 instances from the linear mixed effect model (4.4) with four factors (each has 10 levels) and the true model contains 2-way interactions. The variance parameters are chosen to be $(\sigma_0, \sigma_1, \sigma_2, \sigma) = (2, 1, 0, 0.5)$, or $(2, 1, 0, 1)$, or $(1, 2, 0, 1)$ (fairly strong signal). The empirically estimated λ_1 sometimes can be infinity. In this case, we truncate λ_1 to 5.

The results are shown in the violin plots in Figure 1a to 1c. In all three cases HANOVA with oracle λ achieves Bayes risk and HANOVA with empirical λ performs almost as good as the oracle. The unshrunk linear model suffers from increased noise, as we can see from the first and second plots.

The next simulation suggests HANOVA can fit a high order model without overfitting. In this simulation, the true model is of order 3 but the three way interactions are very weak $((\sigma_0, \sigma_1, \sigma_2, \sigma_3, \sigma) = (2, 1, \mathbf{0.5}, 0, 1))$. The cells can have up to 10 times different weights.

We generate 50 instances and plot the estimation RMSE in Figure 1d. The unshrunk linear model with all the three-way interactions performs poorly, due to overfitting. The second order HANOVA model is slightly better than second order linear model, and the third order model is able to squeeze a little more.

6.2. Real Data.

6.2.1. *IMDb Data.* We ran our algorithm on an IMDb movie ratings data set. This data set contains all the directors who have more than 20 movies on IMDb record. The data set is a $5 \times 20 \times 33$ table with 276 observations (about 8% of all the cells). The three factors in the data set are decade, genre and director.

Since the observed cell is too sparse, we will only fit our hierarchical penalized model to first order and compare it with the unshrunk main effects model. The procedure described in Section 4.2 gives a suggested $\lambda_1 = 0.064$. We further ran cross-validation on the training set and choose to use $\lambda_1 = 0.1$. We use the squared root of the number of movies as cell weights.

The first order HANOVA with λ_1 has test set Root Mean Squared Error (RMSE)= 0.798, while the unshrunk main effects model has RMSE = 0.804 and the cross-validated ridge regression main effects model has RMSE = 0.801.

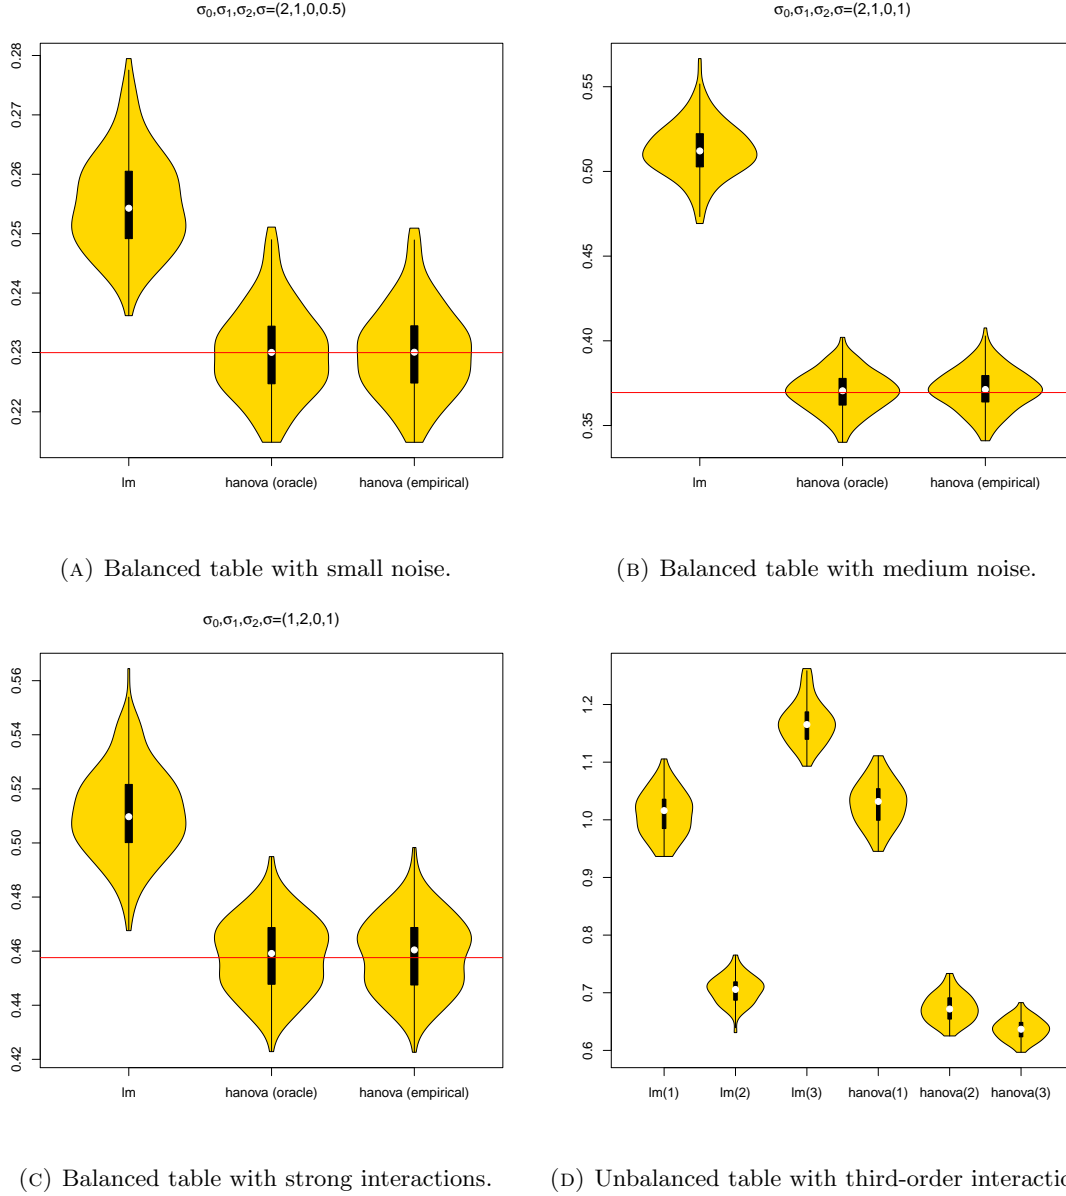
6.2.2. *Zagat Data.* This data set contains survey ratings of 1776 restaurants in New York City in 33 zip-codes, 25 different type of cuisines and 7 price levels. Each rating is from 0 to 30. 20% of the restaurants are left out as test data.

The standard deviation of reviewer random effect is about 7.51. This large random effect makes the ratings of many unpopular restaurants not trustable, as pointed out in Section 1. The number of reviews of each restaurants follow a heavy-tailed distribution. One restaurant has 8922 reviews but 106 restaurants have less than or equal to 5 reviews. The average and median number of reviews are 117.1 and 42.0.

The standard deviation of restaurant random effect (within cell) is about 2.47. This is the lower limit of RMSE of any prediction methods based on the three factors. We use the method described in 5.1 to preprocess the data.

The σ_m estimated by HANOVA (defined in (4.2)) is (0.423, 0.149, 0) (recall the σ^2 is always 1 after the preprocessing procedure 2) and the corresponding $\lambda = (5.71, 45.0, \text{inf})$. This is the case that the data has fairly strong main effects and weak interactions. Since the third-order interactions are too weak, we only fit the first three HANOVA models. The RMSEs are 3.010 (HANOVA grand-mean model), 2.641 (HANOVA main effects model) and 2.638 (HANOVA second-order interactions model). The best λ_2 selected by cross-validation is 72.5 and the corresponding second-order interactions model has RMSE= 2.625. As comparison, the RMSE of random main effects model is 2.640 and the RMSE of price-only random effects model is 2.767.

The RMSE reduction of HANOVA interactions model is more significant in cells that have fewer restaurants. This is one of HANOVA's the main purposes. Among the 355 restaurants in the test data set, 32 of them receive rating adjustment greater than 1 from main effects model to interactions model. Their mean absolute prediction error for these 32 restaurants are reduced from 2.031 to 1.707 by fitting a interactions model by HANOVA.



(A) Balanced table with small noise.

(B) Balanced table with medium noise.

(C) Balanced table with strong interactions.

(D) Unbalanced table with third-order interactions.

FIGURE 1. Comparison of HANOVA with linear model in various simulation settings. The three methods being compared in the first three plots are (left to right) linear model with two-way interactions, HANOVA with oracle λ , and HANOVA with empirically estimated λ . y -axis is the Root Mean Squared Error on all the observed cells. The red horizontal line is the Bayes risk. The methods being compared in the bottom-right plot are linear model (order 1, 2, 3), and HANOVA (order 1, 2, 3) with empirically estimated λ .

REFERENCES

Daniel Barry. Empirical bayes estimation of binomial probabilities in one-way and two-way layouts. *Journal of the Royal Statistical Society. Series D (The Statistician)*, 39(4):pp. 437–453, 1990. ISSN 00390526. URL <http://www.jstor.org/stable/2349087>.

- Rudolf Beran. Asp fits to multi-way layouts. *Annals of Institute of Statistical Mathematics*, 57: 201–220, 2005.
- George Casella. Illustrating empirical bayes methods. *Chemometrics and Intelligent Laboratory Systems*, 16:107–125, 1992.
- Peter J. Diggle, Kung Y. Liang, and Scott L. Zeger. *Analysis of Longitudinal Data*. Oxford University Press, 1994.
- William H. DuMouchel and Jeffrey E. Harris. Bayes Methods for Combining the Results of Cancer Studies in Humans and Other Species. *Journal of the American Statistical Association*, 78(382): 293–308, June 1983. ISSN 01621459. doi: 10.2307/2288631. URL <http://dx.doi.org/10.2307/2288631>.
- Bradley Efron. *Large-Scale Inference: Empirical Bayes Methods for Estimation, Testing, and Prediction*. Cambridge University Press, 2010. ISBN 978-0-521-19249-1.
- Sergey Feldman, Maya R. Gupta, and Bela A. Frigiyk. Multi-task averaging. In *NIPS*, pages 1178–1186, 2012. URL <http://dblp.uni-trier.de/db/conf/nips/nips2012.html/FeldmanGF12>.
- Ronald A. Fisher. The correlation between relatives on the supposition of mendelian inheritance. *Philosophical Transactions of the Royal Society of Edinburgh*, 52:399–433, 1918.
- Andrew Gelman. Analysis of variance - why it is more important than ever. *The Annals of Statistics*, 33(1):1–53, February 2005. ISSN 0090-5364. doi: 10.1214/009053604000001048. URL <http://dx.doi.org/10.1214/009053604000001048>.
- T. J. Hastie and R. J. Tibshirani. *Generalized additive models*. London: Chapman & Hall, 1990. ISBN 0412343908.
- Herbert Robbins. The empirical bayes approach to statistical decision problems. *The Annals of Mathematical Statistics*, 35(1):1–20, 1964. ISSN 00034851. URL <http://www.jstor.org/stable/2238017>.
- Shayle R. Searle. *Linear Models for Unbalanced Data*. Wiley series in probability and mathematical statistics. Applied probability and statistics. J. Wiley, New York, Chichester, 2006. ISBN 978-0-470-04004-1.
- Shayle Robert Searle, George Casella, and Charles E. McCulloch. *Variance components*. Wiley series in probability and mathematical statistics. Applied probability and statistics. J. Wiley, New York, Chichester, 1992. ISBN 0-471-62162-5.
- Alexander Volfovsky and Peter D. Hoff. Hierarchical array priors for anova decompositions. *Annals of Applied Statistics*, 8(1):19–47, 2014.

APPENDIX A. PROOF OF THEOREMS

A.1. **Theorem 1.** First look at

$$(A.1) \quad \beta^{(k)} | \beta^{(k-1)} \sim N(\mathbf{U}_k^T \boldsymbol{\mu}^{(k-1)}, \sigma_{k-1}^2 \mathbf{I}_{d_k}), \quad \boldsymbol{\mu}^{(k-1)} = \mathbf{U}_{k-1} \beta^{(k-1)}$$

The log of the density function is

$$(A.2) \quad \begin{aligned} \log f(\beta^{(k)} | \beta^{(k-1)}) &= -\frac{1}{2} \frac{\|\beta^{(k)} - \mathbf{U}_k^T \boldsymbol{\mu}^{(k-1)}\|^2}{\sigma_{k-1}^2} + \text{const} \\ &= -\frac{1}{2} \frac{\|\mathbf{U}_k(\beta^{(k)} - \mathbf{U}_k^T \boldsymbol{\mu}^{(k-1)})\|^2}{\sigma_{k-1}^2} + \text{const} \\ &= -\frac{1}{2} \frac{\|\boldsymbol{\mu}^{(k)} - \boldsymbol{\mu}^{(k-1)}\|^2}{\sigma_{k-1}^2} + \text{const} \end{aligned}$$

The distribution of \mathbf{y} given β_k is

$$(A.3) \quad \mathbf{y} | \beta_k \sim N(\mathbf{U}_k \beta_k, \sigma^2 \mathbf{I}_n + \sigma_m^2 \mathbf{I}_n + \mathbf{U}_m \mathbf{D} \mathbf{U}_m^T).$$

Here \mathbf{D} is a diagonal matrix

$$(A.4) \quad \mathbf{D} = \begin{pmatrix} (\sigma_{m-1}^2 + \dots + \sigma_k^2) \mathbf{I}_{d_k} & & \\ & \ddots & \\ & & \sigma_{m-1}^2 \mathbf{I}_{d_m - d_{m-1}} \end{pmatrix}.$$

Thus

$$(A.5) \quad \begin{aligned} \log f(\mathbf{y}|\beta_k) &= -\frac{1}{2}(\mathbf{y} - \mathbf{U}_k \beta^{(k)})^T ((\sigma^2 + \sigma_m^2) \mathbf{I}_n + \mathbf{U}_m \mathbf{D} \mathbf{U}_m^T)^{-1} (\mathbf{y} - \mathbf{U}_k \beta^{(k)}) \\ &\quad + \text{const} \\ &= -\frac{1}{2}(\mathbf{y} - \mathbf{U}_k \beta^{(k)})^T \left[\mathbf{U} \begin{pmatrix} (\sigma^2 + \sigma_m^2 + \dots \sigma_k^2) \mathbf{I}_{d_k} & \\ & \tilde{\mathbf{D}} \end{pmatrix} \mathbf{U}^T \right]^{-1} \\ &\quad \cdot (\mathbf{y} - \mathbf{U}_k \beta^{(k)}) + \text{const} \\ &= -\frac{1}{2}(\mathbf{y} - \mathbf{U}_k \beta^{(k)})^T \mathbf{U} \begin{pmatrix} (\sigma^2 + \sigma_m^2 + \dots \sigma_k^2) \mathbf{I}_{d_k} & \\ & \tilde{\mathbf{D}} \end{pmatrix}^{-1} \\ &\quad \cdot \mathbf{U}^T (\mathbf{y} - \mathbf{U}_k \beta^{(k)}) + \text{const} \end{aligned}$$

Notice that $\mathbf{U}^T (\mathbf{y} - \mathbf{U}_k \beta^{(k)}) = \begin{pmatrix} \mathbf{U}_k^T \mathbf{y} - \beta^{(k)} \\ \mathbf{U}_k^{\perp T} \mathbf{y} \end{pmatrix}$, so

$$(A.6) \quad \log f(\mathbf{y}|\beta^{(k)}) = -\frac{\|\beta^{(k)} - \mathbf{U}_k^T \mathbf{y}\|^2}{2(\sigma^2 + \sigma_K^2 + \dots + \sigma_k^2)} - g(\mathbf{y}).$$

The claim is immediately followed by the fact that $\|\beta^{(k)} - \mathbf{U}_k^T \mathbf{y}\|^2 = \|\mathbf{U}_k(\beta^{(k)} - \mathbf{U}_k^T \mathbf{y})\|^2 = \|\boldsymbol{\mu}^{(k)} - \mathbf{P}_k \mathbf{y}\|^2 = \|\boldsymbol{\mu}^{(k)} - \mathbf{y}\|^2 - \|\mathbf{y} - \mathbf{P}_k \mathbf{y}\|^2$. The regularization parameter is

$$(A.7) \quad \lambda_k = \frac{\sigma^2 + \sigma_K^2 + \dots + \sigma_k^2}{\sigma_{k-1}^2}.$$

A.2. Theorem 2. Recall our algorithm is just iteratively maximizing likelihood (4.1) to get $\hat{\beta}_m$ and then find the posterior mean of β given \mathbf{y} and $\beta^{(k)} = \hat{\beta}_m$.

By empirical Bayes estimator of β , I mean just compute the posterior mean of β given \mathbf{y} in (4.4) and plug in whatever σ^2 and σ_k^2 we used to compute λ_k .

The posterior mean of (4.4) can be computed by Tweedie's formula Robbins (1964). The marginal distribution of \mathbf{y} is $N(\mathbf{0}, (\sigma^2 + \sigma_m^2) \mathbf{I}_n + \mathbf{U}_m \boldsymbol{\Sigma} \mathbf{U}_m^T)$, i.e.

$$(A.8) \quad m(\mathbf{y}) \propto \exp\left\{-\frac{1}{2} \mathbf{y}^T ((\sigma^2 + \sigma_m^2) \mathbf{I}_n + \mathbf{U}_m \boldsymbol{\Sigma} \mathbf{U}_m^T)^{-1} \mathbf{y}\right\}$$

so the posterior mean is given by

$$(A.9) \quad \begin{aligned} E[\boldsymbol{\mu}|\mathbf{y}] &= \mathbf{y} + \sigma^2 \nabla \log m(\mathbf{y}) \\ &= \mathbf{y} - \sigma^2 ((\sigma^2 + \sigma_m^2) \mathbf{I}_n + \mathbf{U}_m \boldsymbol{\Sigma} \mathbf{U}_m^T)^{-1} \mathbf{y} \\ &= \frac{\sigma_m^2}{\sigma^2 + \sigma_m^2} \mathbf{y} \\ &\quad + \mathbf{U} \left(\frac{\sigma^2}{\sigma^2 + \sigma_m^2} \mathbf{I}_n - \sigma^2 \begin{pmatrix} (\sigma^2 + \sigma_m^2) \mathbf{I}_{d_m} + \boldsymbol{\Sigma} & \\ & (\sigma^2 + \sigma_m^2) \mathbf{I}_{n-d_m} \end{pmatrix}^{-1} \right) \mathbf{U}^T \mathbf{y} \\ &= \frac{\sigma_m^2}{\sigma^2 + \sigma_m^2} \mathbf{y} + \sum_{k=1}^m \left(\frac{\sigma^2}{\sigma^2 + \sigma_m^2} - \frac{\sigma^2}{\sigma^2 + \sigma_m^2 + \dots + \sigma_{k-1}^2} \right) \mathbf{V}_k \mathbf{V}_k^T \mathbf{y} \end{aligned}$$

On the other hand, in the case of a balanced table, each step of our algorithm produces

$$\begin{aligned}
 \hat{\boldsymbol{\mu}}^{(k)} &= \mathbf{U}_k \mathbf{U}_k^T \left(\frac{1}{1 + \lambda_k} \mathbf{y} + \frac{\lambda_k}{1 + \lambda_k} \hat{\boldsymbol{\mu}}^{(k-1)} \right) \\
 (A.10) \quad &= \mathbf{U}_k \mathbf{U}_k^T \left(\frac{\sigma_{k-1}^2}{\sigma^2 + \sigma_m^2 + \dots + \sigma_{k-1}^2} \mathbf{y} + \frac{\sigma^2 + \sigma_m^2 + \dots + \sigma_k^2}{\sigma^2 + \sigma_m^2 + \dots + \sigma_{k-1}^2} \hat{\boldsymbol{\mu}}^{(k-1)} \right), \\
 &\quad k = 1, \dots, m
 \end{aligned}$$

and finally

$$(A.11) \quad \hat{\boldsymbol{\mu}} = \mathbb{E}[\boldsymbol{\mu} | \mathbf{y}, \hat{\mathbf{y}}^{(k)}] = \frac{\sigma_m^2}{\sigma^2 + \sigma_m^2} \mathbf{y} + \frac{\sigma^2}{\sigma^2 + \sigma_m^2} \hat{\mathbf{y}}^{(m)}$$

It suffices to show (A.9) and (A.11) are actually the same. Notice that $\mathbf{U}_k \mathbf{U}_k^T \hat{\boldsymbol{\mu}}^{(j)} = \hat{\boldsymbol{\mu}}^{(j)}$ for $j < k$ because $\hat{\boldsymbol{\mu}}^{(j)} \in \mathcal{S}_j \subset \mathcal{S}_k$ and $\mathbf{U}_k \mathbf{U}_k^T$ is just the projection matrix onto \mathcal{S}_k . With some calculation,

$$\begin{aligned}
 \hat{\boldsymbol{\mu}}^{(k)} &= \sum_{k=1}^m \frac{\sigma^2 + \sigma_m^2}{\sigma^2 + \sigma_m^2 + \dots + \sigma_k^2} \frac{\sigma_{k-1}^2}{\sigma^2 + \sigma_m^2 + \dots + \sigma_{k-1}^2} \mathbf{U}_k \mathbf{U}_k^T \mathbf{y} \\
 (A.12) \quad &= (\sigma^2 + \sigma_m^2) \sum_{k=1}^m \left(\frac{1}{\sigma^2 + \sigma_m^2 + \dots + \sigma_k^2} - \frac{1}{\sigma^2 + \sigma_m^2 + \dots + \sigma_{k-1}^2} \right) \mathbf{U}_k \mathbf{U}_k^T \mathbf{y}
 \end{aligned}$$

Since $\mathbf{U}_k \mathbf{U}_k^T \mathbf{y} = \sum_{j=1}^k \mathbf{V}_j \mathbf{V}_j^T \mathbf{y}$, it is easy to see that

$$(A.13) \quad \frac{\sigma^2}{\sigma^2 + \sigma_m^2} \hat{\boldsymbol{\mu}}^{(k)} = \sigma^2 \sum_{k=1}^m \left(\frac{1}{\sigma^2 + \sigma_m^2} - \frac{1}{\sigma^2 + \sigma_m^2 + \dots + \sigma_{k-1}^2} \right) \mathbf{V}_k \mathbf{V}_k^T \mathbf{y}$$

The theorem is immediately proved if we plug this in (A.11) and compare it to (A.9).